CLAIMS

What is claimed is:

1. A storage system comprising a redundant array of multicast storage areas.

2. The storage system of claim 1, wherein:
   the multicast devices are adapted to communicate across a network via encapsulated packets which are split-ID packets comprising both an encapsulating packet and an encapsulated packet; and
   each of any split-ID packets also includes an identifier that is split such that a portion of the identifier is obtained from the encapsulated packet while another portion is obtained from a header portion of the encapsulating packet.

3. The storage system of claim 1, wherein the storage areas of the redundant array share a common multicast address.

4. The storage system of claim 1, comprising a plurality of RAID sets wherein each raid set comprises a plurality of storage areas sharing a common multicast address.

5. A network comprising a first device and a plurality of storage devices wherein the first device stores a unit of data on each of the storage devices via a single multicast packet.

6. A network of multicast devices which disaggregate at least one RAID function across multiple multicast addressable storage areas.

7. The network of claim 6 wherein the at least one RAID function is also disaggregated across multiple device controllers.

8. A storage system comprising a redundant array of multicast storage areas wherein the system supports auto-annihilation of mooted read requests.

9. The system of claim 8 wherein auto-annihilation comprises the first device responding to a read request commanding other devices to disregard the same read request.

10. The system of claim 9 wherein auto-annihilation comprises a device that received a read request disregarding the read request if a response to the read request from another device is detected.

11. A storage system comprising a dynamic mirror.

12. The storage system of claim 11 wherein the dynamic mirror includes a mirrored storage area and at least one corresponding map of incomplete writes.

13. The storage system of claim 11 wherein the dynamic mirror comprises N storage devices and M maps of incomplete writes where M is at least 1 and at most 2*N.

14. The storage system of claim 13 wherein the map comprises a set of entries wherein each entry is either an LBA or a hash of an LBA of a storage block of a storage area being mirrored.

15. The system of claim 13 comprising at least one process monitoring storage area ACKs sent in response to write commands, the process updating any map associated with a particular area whenever a write command applicable to the area is issued, the process also sending an ACK on behalf of any storage area for which the process did not detect an ACK.

16. The system of claim 55 wherein updating a map comprises setting a flag whenever an ACK is not received and clearing a flag whenever an ACK is received.